

Amendments to the Claims

1 Claim 1 (currently amended): A method of improving intrusion detection in a computing
2 network, comprising steps of:

3 defining a plurality of intrusion suspicion levels for use when performing intrusion
4 detection processing on inbound communications destined for a computing device on the
5 computing network;

6 for each of a plurality of potential intrusion events, defining a set of at least one
7 conditions which describe the potential intrusion event;

8 associating one of the defined intrusion suspicion levels with each of the sets of
9 conditions;

10 defining a plurality of sensitivity levels for filtering intrusion events when performing the
11 intrusion detection processing; and

12 performing intrusion detection for a particular inbound communication received for the
13 computing device, further comprising steps of:

14 determining whether any of the at least one sets of conditions are matched; and
15 if so, using a currently-applicable one of the defined sensitivity levels, in concert
16 with the defined intrusion suspicion levels level associated with the matched conditions, to
17 determine if [[a]] the particular inbound communication destined for the computing device
18 should be treated as an intrusion event.

Claim 2 (canceled)

Serial No. 10/058,689

-2-

RSW920020011US1

1. Claim 3 (currently amended): The method according to Claim [[2]] 1, wherein the determining
2 step further comprises comparing current conditions in the computing device to predetermined
3 conditions which signal a potential intrusion the conditions defined in at least one of the sets.

1. Claim 4 (currently amended): The method according to Claim 3, wherein the current conditions
2 in the computing device comprise contents of the particular inbound communication.

1. Claim 5 (currently amended): The method according to Claim 4, wherein the current conditions
2 in the computing device further comprise a protocol state of a protocol stack which processes the
3 particular inbound communication.

1. Claim 6 (currently amended): The method according to Claim 1, further comprising the step of
2 taking one or more defensive actions when the using step determines upon determining that the
3 particular inbound communication should be treated as an intrusion event.

1. Claim 7 (original): The method according to Claim 6, wherein the defensive actions are
2 determined by consulting intrusion detection policy information.

1. Claim 8 (currently amended): The method according to Claim [[6]] 7, wherein the intrusion
2 detection policy information is stored in a network-accessible repository.

1. Claim 9 (currently amended): The method according to Claim 1, wherein the using step further

2 comprises comparing the particular inbound communication to defined at least one set of
3 conditions represents one or more attack signatures.

1 Claim 10 (original): The method according to Claim 9, wherein at least one of the attack
2 signatures is a class signature representing a class of attacks.

1 Claim 11 (currently amended): The method according to Claim [[9]] 1, wherein each of the at
2 least one set of conditions is attack signatures are specified as conditions as a condition part in an
3 intrusion detection [[rules]] rule, and wherein each of the intrusion detection rules further
4 specifies at least one action comprises one or more actions that are to be taken upon determining
5 when the using step determines that the particular inbound communication should be treated as
6 an intrusion event.

1 Claim 12 (currently amended): The method according to Claim 1, wherein the performing
2 [[using]] step operates in the computing device for which the particular inbound communication
3 is destined.

1 Claim 13 (currently amended): The method according to Claim 12, wherein the performing
2 [[using]] step operates within layer-specific intrusion detection logic executing in a protocol
3 stack running on the computing device.

1 Claim 14 (currently amended): The method according to Claim 1, wherein the performing

2 [[using]] step operates in a network device which analyzes communications directed to the
3 computing device for which the particular inbound communication is destined.

1 Claim 15 (currently amended): The method according to Claim 1, further comprising steps of:
2 ————— for each of a plurality of potential intrusion events, defining a set of one or more
3 conditions which describe the potential intrusion event;
4 ————— associating a sensitivity level with each of the sets of conditions; and
5 ————— determining a suspicion level of the particular inbound communication;
6 ————— wherein the using step further comprises consulting a stored mapping between each of the
7 defined sensitivity levels and each of the defined intrusion suspicion levels, using the currently-
8 applicable one of the defined sensitivity levels and the intrusion suspicion level associated with
9 the matched conditions, to determine if determines that the particular inbound communication
10 should be treated as an intrusion event when conditions pertaining to the particular inbound
11 communication match a selected one of the sets of conditions and the determined suspicion level
12 maps to the sensitivity level associated with the selected set of conditions.

Claims 16 - 21 (canceled)

1 Claim 22 (currently amended): A system for improving intrusion detection in a computing
2 network, comprising:
3 means for defining a plurality of intrusion suspicion levels defined for use when
4 performing intrusion detection processing on inbound communications destined for a computing

5 device on the computing network;

6 for each of a plurality of potential intrusion events, a defined set of at least one conditions

7 which describe the potential intrusion event;

8 means for associating one of the defined intrusion suspicion levels with each of the

9 defined sets of conditions;

10 a plurality of sensitivity levels defined for filtering intrusion events when performing the

11 intrusion detection processing; and

12 means for performing intrusion detection for a particular inbound communication

13 received for the computing device, further comprising:

14 means for determining whether any of the at least one defined sets of conditions

15 are matched; and

16 if so, means for using a currently-applicable one of the defined sensitivity levels,

17 in concert with the defined intrusion suspicion levels level associated with the matched

18 conditions, to determine if [[a]] the particular inbound communication destined for the

19 computing device should be treated as an intrusion event.

Claim 23 (canceled)

1 Claim 24 (currently amended): The system according to Claim [[23]] 22, wherein the means for

2 determining further comprises means for comparing current conditions in the computing device

3 to predetermined conditions which signal a potential intrusion the conditions defined in at least

4 one of the sets.

1 Claim 25 (currently amended): The system according to Claim 22, further comprising means for
2 taking one or more defensive actions when the means for using determines upon determining that
3 the particular inbound communication should be treated as an intrusion event, wherein the
4 defensive actions are determined by consulting intrusion detection policy information.

1 Claim 26 (currently amended): The system according to Claim 22, wherein each of the means
2 for using further comprises means for comparing the particular inbound communication to at
3 least one set of conditions is one or more attack signatures, wherein the attack signatures are
4 specified as conditions a condition part in an intrusion detection rules rule, and wherein each of
5 the intrusion detection rules further comprises specifies at least one action one or more actions
6 that are to be taken upon determining when the means for using determines that the particular
7 inbound communication should be treated as an intrusion event.

1 Claim 27 (currently amended): The system according to Claim 22, further comprising:
2 — for each of a plurality of potential intrusion events, means for defining a set of one or
3 more conditions which describe the potential intrusion event;
4 — means for associating a sensitivity level with each of the sets of conditions; and
5 — means for determining a suspicion level of the particular inbound communication;
6 — wherein the means for using further comprises means for consulting a stored mapping
7 between each of the defined sensitivity levels and each of the defined intrusion suspicion levels,
8 using the currently-applicable one of the defined sensitivity levels and the intrusion suspicion

9 level associated with the matched conditions, to determine if determines that the particular
10 inbound communication should be treated as an intrusion event when conditions pertaining to the
11 particular inbound communication match a selected one of the sets of conditions and the
12 determined suspicion level maps to the sensitivity level associated with the selected set of
13 conditions.

Claims 28 - 31 (canceled)

1 Claim 32 (currently amended): A computer program product for improving intrusion detection
2 in a computing network, the computer program product embodied on one or more computer-
3 readable media and comprising:
4 computer-readable program code means for defining a plurality of intrusion suspicion
5 levels for use when performing intrusion detection processing on inbound communications
6 destined for a computing device on the computing network;
7 for each of a plurality of potential intrusion events, computer-readable program code
8 defining a set of at least one conditions which describe the potential intrusion event;
9 computer-readable program code associating one of the defined intrusion suspicion levels
10 with each of the sets of conditions;
11 computer-readable program code defining a plurality of sensitivity levels for filtering
12 intrusion events when performing the intrusion detection processing; and
13 computer-readable program code for performing intrusion detection for a particular
14 inbound communication received for the computing device, further comprising:

15 computer-readable program code for determining whether any of the at least one
16 sets of conditions are matched; and
17 if so, computer-readable program code [[means]] for using a currently-applicable
18 one of the defined sensitivity levels, in concert with the defined intrusion suspicion levels level
19 associated with the matched conditions, to determine if [[a]] the particular inbound
20 communication destined for the computing device should be treated as an intrusion event.

Claim 33 (canceled)

1 Claim 34 (currently amended): The computer program product according to Claim [[33]] 32,
2 wherein the computer-readable program code [[means]] for determining further comprises
3 computer-readable program code [[means]] for comparing current conditions in the computing
4 device to predetermined conditions which signal a potential intrusion the conditions defined in at
5 least one of the sets, the current conditions in the computing device comprising contents of the
6 particular inbound communication.

1 Claim 35 (currently amended): The computer program product according to Claim [[33]] 32,
2 wherein the computer-readable program code [[means]] for determining further comprises
3 computer-readable program code [[means]] for comparing current conditions in the computing
4 device to predetermined conditions which signal a potential intrusion the conditions defined in at
5 least one of the sets, the current conditions in the computing device comprising contents of the
6 particular inbound communication and a protocol state of a protocol stack which processes the

7 particular inbound communication.

1 Claim 36 (currently amended): The computer program product according to Claim 32, further
2 comprising computer-readable program code [[means]] for taking one or more defensive actions
3 upon determining when the computer-readable program code means for using determines that the
4 particular inbound communication should be treated as an intrusion event, wherein the defensive
5 actions are determined by consulting intrusion detection policy information stored in a policy
6 repository.

1 Claim 37 (currently amended): The computer program product according to Claim [[1]] 32,
2 wherein the computer-readable program code means for using further comprises computer-
3 readable program code means for comparing the particular inbound communication to defined at
4 least one set of conditions represents one or more attack signatures, wherein at least one of the
5 attack signatures is a class signature representing a class of attacks.

1 Claim 38 (currently amended): The computer program product according to Claim 32, wherein
2 the computer-readable program code [[means]] for [[using]] performing operates in the
3 computing device for which the particular inbound communication is destined.

1 Claim 39 (currently amended): The computer program product according to Claim 32, wherein
2 the computer-readable program code [[means]] for [[using]] performing operates in a network
3 device which analyzes communications directed to the computing device for which the particular

4 inbound communication is destined.

1 Claim 40 (currently amended): The computer program product according to Claim 32, further
2 comprising:

3 —— computer-readable program code means for specifying, for each of a plurality of potential
4 intrusion events, a set of one or more conditions which describe the potential intrusion event;
5 —— computer-readable program code means for associating a sensitivity level with each of the
6 sets of conditions; and

7 —— computer-readable program code means for determining a suspicion level of the
8 particular inbound communication;

9 —— wherein the computer-readable program code [[means]] for using further comprises
10 computer-readable code for consulting a stored mapping between each of the defined sensitivity
11 levels and each of the defined intrusion suspicion levels, using the currently-applicable one of the
12 defined sensitivity levels and the intrusion suspicion level associated with the matched
13 conditions, to determine if determines that the particular inbound communication should be
14 treated as an intrusion event when conditions pertaining to the particular inbound communication
15 match a selected one of the sets of conditions and the determined suspicion level maps to the
16 sensitivity level associated with the selected set of conditions.

Claims 41 - 44 (canceled)

1 Claim 45 (new): The method according to Claim 6, wherein the defensive actions are specified

Serial No. 10/058,689

-11-

RSW920020011US1

2 as actions in a rule in which the matched conditions are specified.

1 Claim 46 (new): The method according to Claim 6, wherein at least one of the defensive actions
2 comprises discarding the particular inbound communication.

1 Claim 47 (new): The method according to Claim 6, wherein at least one of the defensive actions
2 comprises limiting at least one of resources or traffic associated with a connection on which the
3 particular inbound communication is received.

1 Claim 48 (new): The method according to Claim 6, wherein at least one of the defensive actions
2 comprises dynamically dropping a deny filter into the computing network to shun subsequent
3 traffic.

1 Claim 49 (new): The method according to Claim 6, wherein at least one of the defensive actions
2 comprises reporting the intrusion event to one or more entities.

1 Claim 50 (new): The method according to Claim 49, wherein reporting the intrusion event to
2 one or more entities further comprises sending an alert to a management component external
3 from the computing device for which the particular inbound communication is destined.

1 Claim 51 (new): The method according to Claim 49, wherein reporting the intrusion event to
2 one or more entities further comprises writing at least one event record to at least one of a system

3 log and a console.

1 **Claim 52 (new):** The method according to **Claim 49**, wherein reporting the intrusion event to
2 one or more entities further comprises recording inbound communications associated with the
3 intrusion event in at least one of a trace or other repository.

1 **Claim 53 (new):** The method according to **Claim 49**, wherein reporting the intrusion event to
2 one or more entities further comprises writing statistics records on normal behavior to establish
3 baselines as to what constitutes abnormal behavior for the inbound communications.

1 **Claim 54 (new):** The method according to **Claim 1**, wherein at least one of the defined sets of
2 conditions specifies a current system state of the computing device.

1 **Claim 55 (new):** The method according to **Claim 1**, wherein at least one of the defined sets of
2 conditions specifies at least one threshold reached at the computing device.

1 **Claim 56 (new):** The method according to **Claim 1**, wherein at least one of the defined sets of
2 conditions specifies at least one state transition to be caused, at the computing device, upon
3 receiving the particular inbound communication.

1 **Claim 57 (new):** The method according to **Claim 1**, wherein the currently-applicable sensitivity
2 level is specified, for the computing device, by a systems administrator.

1 Claim 58 (new): The method according to Claim 1, wherein the currently-applicable sensitivity
2 level is specified, for the computing device, by configuration data in a stored repository.